STATEMENT OF WORK
TIME SCALE MEASUREMENT SYSTEM

INTRODUCTION

The Time Scale Measurement System at the Boulder facility of the National Institute of Standards and Technology (NIST) is well over 20 years old. Though it has performed very well, there is a need for improved stability, and it is becoming increasingly hard to obtain replacement parts. Therefore NIST is initiating a program to improve and modernize this equipment.

Present System

The current time scale measurement system works by measuring the phase, or time difference, between the 5 MHz outputs of highly stable atomic frequency standards. To provide improved accuracy and stability of the measurement system, a low noise frequency synthesizer, which is referenced to one of the atomic frequency standards, generates an output signal that is coherently offset by 10 Hz, (i.e. it generates 4,999,990 Hz). The nominal 10 Hz beat frequencies between the output of the synthesizer and the 5 MHz outputs of each of a number of other atomic frequency standards are the signals that are actually measured. These low frequency beat signals provide enormously leveraged measurements. For example, a 1 Hz offset between frequency standards at 5 MHz, (20 nanoseconds per day in phase) will appear also as a 1 Hz offset at 10 Hz. Thus a measurement requiring 2×10^{-13} resolution at 5 MHz becomes a measurement requiring only 1×10^{-7} resolution at 10 Hz.

Each frequency standard requires its own measurement channel. In each channel there is a counter that continuously counts the number of 10 Hz cycles between the offset reference and the clock under test. Each time a measurement cycle is initiated, a second register measures the phase difference (time interval) between a zero crossing of the 10Hz beat frequency of the reference standard and the next zero crossing of the standard being measured. Specifically, this phase difference counter counts the number of cycles of an internal 10 MHz source in the interval from a positive going zero crossing of the reference 10 Hz signal to the next positive going zero crossing of the test channel. The full-scale range of this phase difference measurement is 1 million counts; because of the leverage provided by the down-conversion from 5 MHz to 10 Hz, 1 count corresponds to a time difference of 0.2 ps between the two input frequencies. (One cycle of the 10 Hz beat frequency corresponds to one cycle at 5 MHz, and therefore 200 ns). Figure 1 shows a block diagram of the existing system.

REQUIREMENTS FOR NEW SYSTEM

The new system is expected to have the same general layout as the existing system as shown in Fig. 1. The details need not be the same but functionally it should be similar. (It is NIST's technical judgment that designs based on measurements at 1 pulse per second cannot satisfy the

requirements herein, and all such designs are unacceptable.). The new system should provide a method for translating the input frequency (at 5 and 100 MHz, and possibly 10 MHz) to a lower frequency of order 10 Hz, combined with a method of making phase-difference measurements at this lower intermediate frequency. In addition, the system shall include ancillary hardware to interface the measurement system to a computer, a rack-mounted chassis, a power supply, etc. The detailed requirements are set out more fully below.

The Contractor must have demonstrated knowledge and experience with the design, fabrication, and testing of low phase noise, high frequency stability measurement equipment.

Input Signals

The new system shall have channels that accept input frequencies from hydrogen masers or high stability cesium beam frequency standards at 5 and 100 MHz (10 MHz optional). The generated beat frequencies shall be approximately 10 Hz in all cases. The exact intermediate frequency may be changed following discussions with the NIST Contracting Officer's Technical Representative. In general, a lower intermediate frequency will increase the leverage in the measurement process, but may also increase the 1/f noise in the hardware and will also decrease the maximum frequency offset that the system can handle. The final design shall explain the choice of intermediate frequency including an evaluation that includes (but need not be limited to) a discussion of these considerations. In the following discussion, the term "10 Hz beat frequency" is used with the understanding that this is a descriptive term, and that the exact frequency that is used may differ based on these considerations. Each channel may be designed for only one specific input frequency, although a method for reconfiguring a channel at any time (using some combination of jumpers, switches or software commands) is desirable. The channels shall be designed to work with RF input amplitudes of 1+/-0.1 volts rms. The input impedance shall be 50 ohms with a maximum return loss of 30 dB at 5MHz and 20 dB at 100MHz. All RF inputs shall have DC isolated grounds and the connectors shall be SMA. When the input sine waves are used to trigger a digital circuit, the trigger point shall be +1 volt with a positive slope.

NIST will supply a 5 MHz signal to be used as the reference clock. The measurement system shall generate the 100 MHz reference frequency (and 10 MHz optional frequency) internally. The system shall use these signals as the reference clocks from which the 10 Hz offset frequencies are synthesized. It is highly desirable, but not mandatory, that the system also accept a 100 MHz external reference frequency instead of the 5 MHz signal, and that the 5 MHz and 10 MHz signals be generated internally. These configuration changes need not be automatic but can be set up manually. However, a change in reference frequency or reference clock must be able to be made without having to power down the measurement system.

Triggered Individual Measurements. A trigger signal to initiate the measurement cycle will be provided by a NIST supplied computer through a standard computer interface. This trigger signal may be implemented by the Contractor using any i/o command that is native to the host operating

system, such as "write" or "i/o control," and this single command shall initiate whatever discrete sub-operations are required to perform the complete measurement. (See Computer Interface below.) This trigger signal shall initiate a cycle that measures the time difference between the reference clock and all of the other clocks that have been declared as active. The cycle shall begin on the next zero-crossing of the beat frequency. NIST will design its software so that the command to start a cycle will not be given so close to the zero crossing of the beat frequency that a potential race condition may be generated. The current NIST design guarantees that the start cycle signal is given no later than 15 ms before the next zero crossing of the beat frequency and no sooner than 15 ms after the last zero crossing. NIST will support this requirement in the new system. A larger delay gap will be considered provided it is adequately justified by the Contractor. NIST will consider designs in which the actual measurements start some time after the command is given, provided that this delay is the same for all channels and is a digitally-derived integer multiple of one half of the beat period. (In the following discussion, the phrase "next cycle" is understood to include a system with a start delay that satisfies this requirement.) Systems in which the clocks are measured at different times (using multiplexed channels, for example), or in which the actual start of the measurements are not digitally derived from the beat frequency in the same way on each cycle are not acceptable.

Automatic Periodic Measurements (desirable). Although the current hardware does not support this capability, it would be useful, but not mandatory, if the measurement hardware could be configured so that it automatically initiated a measurement cycle periodically, where the period between cycles was derived by counting the 10 Hz beat frequency of the reference clock. For example, the controlling software would send a command to the measurement system to initiate a measurement cycle on every Nth count of the beat frequency of the reference clock. The measurement cycles would then be performed automatically until a stop command was given. The host operating system would be responsible for reading the results fast enough so that data were not lost; an optional "data overrun" flag might be provided to show that this had not happened. Even at the shortest possible period, the data rate is not very high, and keeping up with it should not present much of a problem. While this capability probably would not used in the time scale, it would be useful for characterizing the frequency stability of an oscillator.

Typically, the measurement cycle shall consist of two actions in each clock measurement channel that has been declared to be active by the controlling software (see below). First, the next zero crossing of the 10 Hz signal from the reference clock shall cause each interface to strobe its accumulated number of integer counts of the beat frequency to a holding register. Second, this pulse shall initialize and then start a counter that counts a reference frequency. This counter is stopped by the next zero crossing of the 10 Hz signal from the clock being measured. (The current system uses a reference frequency of 10 MHz for this purpose. A higher reference frequency would provide additional resolution and is desirable if this increased resolution can be supported by the rest of the system.) Other means of measuring the fractional-cycle phase difference will be considered, provided they are fully digital in nature.

In addition, each device shall contain a number of status and error-checking signals as described below. All of these data shall be stored in static registers that are not over-written until the next measurement cycle is initiated by the controlling software (or by the automatic periodic scheduler, if this capability is supported).

Number of Input Channels

A minimum of 24 input channels in each system shall be required. There shall be a minimum of 8 channels for 100 MHz and 16 channels for 5 MHz. Depending on price more channels for 10 MHz may be added or replace some of the 5 MHz channels. Ultimately 192 channels may be required.

Output Information

The information outputted from each measurement channel shall be: (1) the number of cycles of the 10 Hz beat frequency. (2) The phase difference between the 10 Hz beat from the reference oscillator and the test oscillators. (3) Error and status information as described below.

Even under worst-case circumstances, a measurement cycle shall be completed within 2 periods of the beat frequency after it has been initiated (that is, within 200 ms in the current system). The hardware (and any drivers or interface software supplied by the Contractor) shall be capable of supporting "back-to-back" cycles indefinitely, assuming that the NIST-supplied computer has sufficient resources to support this rate. Since the measurements are initiated by the NIST control software, there should be nothing in the Contractor-supplied system that limits or controls the maximum or minimum interval between cycles (beyond the maximum repetition rate limit described in the previous sentence). This requirement may be relaxed somewhat in the case that some part of the measurement cycle encounters a time-out error, but the system shall support measurements made at a rate of at least once per second even if one clock module has continuous time-out errors. If automatic periodic measurement cycles are supported, an error in one of the cycles may terminate the sequence, and the need to support continuous measurements may be waived in the case of a hardware error.

Computer Interface

The computer interface shall be standard computer hardware and shall use an industry-standard interface design such as the PCI bus or the VME bus. Designs based on the (E)ISA standards are not acceptable, since these interfaces are becoming obsolete now and are unlikely to be widely supported in a few years. The controlling computer and its operating system will be supplied by NIST, and will be chosen based on the detailed interface design proposed by the contractor. NIST will consider a design based on any generally-available operating system that is currently supported by its supplier. Examples of acceptable operating systems include (but are not limited to) Microsoft Windows NT/2000, Compaq/Digital Tru64 UNIX, SunOS, Linux and FreeBSD. (Note that these operating systems only support certain hardware configurations, so that the computer hardware and its operating system cannot be specified independently of each other.) Operating systems or

computer hardware configurations that are unique to the Contractor or which are not available on the general open market are not acceptable. In addition to the hardware needed to meet the NIST requirements, the Contractor shall supply a driver program (and library routines, if necessary) that allows programs written in standard high-level languages (e.g., C or FORTRAN) to communicate with the hardware and to control all of its features using both interrupt-driven and polled modes. These components of the system may be supplied by third-parties, provided that the Contractor assumes overall system-level responsibility for its operation.

As a minimum, the driver and the measurement hardware shall support the following capabilities. In all cases, the command syntax shall contain a parameter that specifies which measurement module is being addressed, and this parameter shall have sufficient dynamic range to support the maximum number of clocks that can be measured by the hardware.

1. *The counter and the holding register for the beat frequency*. This register shall be large enough to support counting the chosen beat frequency continuously for at least 1 year (3×10^7 s). This requirement can probably be provided using a register that is 32-bits wide (with the overflow bit in the most significant position), but a smaller register is acceptable if it can satisfy this requirement. The counter shall contain an overflow bit that is set when the maximum count is reached *and is only reset by a special software command.* The counter and the holding register shall support the following commands: reset counter to zero, reset overflow bit to zero, read holding register and overflow bit in one operation, read overflow only. The read operations shall not alter the contents of the registers.
2. *The phase difference register*. This register shall be large enough to support counting the chosen reference frequency for a full cycle of the beat period. A time-out flag shall be provided if the count in this register exceeds the maximum possible value for the design. This flag may be part of the counter and may be triggered on an even power of 2. For example, the maximum count in the current system is 10^6 (counting 10 MHz for 0.1 s). This maximum count can be held in 20 bits, so that the 21st bit of the register can be used to signal an overflow time-out. The counter shall stop counting when an overflow condition is sensed and the hardware shall then report that the measurement cycle has been completed. The software will detect the time out error by noting that the overflow bit is set on the read. A separate error flag is not required. A design in which this counter rolls over through 0 or does not stop is unacceptable. In all cases the stored data shall be readable under software control and shall not be over-written until the next measurement cycle is initiated.
3. *Enable measurements, disable measurements*. These commands shall initialize the hardware for a specific measurement channel, shall start and stop the beat frequency counters and shall enable or disable interrupts from this channel.

4. *Status flags*. In addition to the registers described above, each clock module shall support the following status flags: 10 Hz beat frequency present, reference frequency for counter present, start measurement pulse received, and measurement completed.

It is also desirable to display the state of at least some of these status bits using front-panel LEDs. These displays may be limited based on the size of the front panel and the free space left on it after all of the required connectors have been positioned. The driver and any software libraries shall interface the measurement system to the computer using the protocols and conventions specified by the provider of the operating system for native drivers. (Drivers based only on commands such as "peek" and "poke" are not acceptable.) When used in polled mode, the clock modules shall be addressable using the standard read and write commands of the operating system (or using array references if the driver supports memory mapping). All such operations shall be non-blocking by default, if possible. That is, they should return immediately with a status equivalent to "in progress" or "would block" if the operation is not yet completed. (If they are blocking by default, a parameter must be provided to configure them as non-blocking.) A system which supports only blocking i/o operations is unacceptable.

When used in interrupt-driven mode, the software shall transmit these interrupts to the application using whatever method is standard for that operating system: ("signals" or "messages" or "event flags" or …). The system shall be designed with one or two interrupts, with each interrupt providing additional information via a status register. For example, all errors may be vectored to a single interrupt, with a separate status register providing additional details. Likewise, a single interrupt may be used to signal that any operation has completed normally, with additional details provided by ancillary registers. The data in these registers may be returned as part of the interrupt service itself, or reading them may be accomplished by a separate i/o operation. The hardware and software shall support a mask register so that any combination of interrupts may be either accepted or ignored. This register shall include the capability of turning interrupts off altogether and running the hardware in polled mode.

Modular Design

The design shall be modular in nature so that more measurement channels can be added at a later time. In addition, the system shall be designed so that there are a minimum number of single points of failure. As much as possible, each clock module should be electrically independent of all of the others, so that a failure of one module does not stop the entire system. (Providing separate power regulators on each clock module may be desirable both for this reason and to improve the immunity of the overall system to cross-talk noise.) Similar considerations should govern the software: to the greatest extent possible, every command should support an explicit time-out facility so that a failure of one register or module does not cause the whole system to get stuck waiting for a response that will never come. (This requirement is in addition to providing non-blocking i/o commands.)

Hardware

At least 24 channels shall fit in a standard 19-inch rack with a maximum height of 34 cm (14 inches) and a maximum depth of 53 cm (21 inches). The unit shall run off of 120 VAC, single phase in a normal room temperature environment. Provision for an external DC back up power source (24 volts) is highly desirable but not mandatory. The system shall be designed so that a single source of the 10 Hz offset frequency, and any necessary trigger signals, can be distributed to several racks. Designs based on standard, generally-available components are strongly encouraged; the use of components that are only available from a single source is discouraged, but will be accepted if no other choice will satisfy the requirements.

NIST will consider configurations in which the computer is built into the same rack as the measurement system as well as systems in which the computer and the measurement system are in separate chassis connected by an interface cable (the current arrangement). While a single chassis configuration has some advantages with respect to power and overall size, these advantages may not be substituted for the stability requirements outlined below. The stability requirement is much more important, and NIST is prepared to sacrifice size and power consumption to improve it.

Stability

The maximum acceptable instability for a time difference measurement on the 100 MHz channels is 0.06 ps TDEV at 1 s for a measurement rate of once per second. The corresponding maximum Allan deviation at 1s is $\sigma_y(\tau=1s) = 1x10^{-13}$. For longer time intervals, τ, the Allan deviation shall not exceed that shown in Fig. 2 for the 100 MHz stability requirement. Beyond $\tau = 1x10^{6}$ s the Allan deviation shall not exceed $1x10^{-18}$. Better performance is highly desirable and a goal for the 100 MHz signal is $\sigma_y(\tau=1s) = 1x10^{-14}$. The stability goal should be such that the Allan deviation decreases as $1/\tau$, or faster, to a level of $1x10^{-19}$ and never exceeds this value as shown in Fig. 2. The Allan deviation of the measurement systems shall be determined by using two independent systems to measure the same clock, and it shall be assumed that the two systems are equivalent. Similar tests shall also be performed with two channels on the same system.

The stability requirement for the 5 MHz signal is also shown in Fig. 2. Beyond $\tau = 1x10^{7}$ s the Allan deviation shall not exceed $1x10^{-18}$. The stability requirement for the optional 10 MHz input is a factor of 2 lower than that for the 5 MHz signal. The 5 MHz goal is a noise level no more than 10 times that of the 100 MHz goal. The 10 MHz goal is no more than five times higher than the 100 MHz goal.

For the 100 MHz signal the temperature sensitivity of the measurement hardware shall be less than 2 ps per C in the range of 20 to 30 C ambient room temperature. For the 5 MHz signal the temperature sensitivity of the measurement hardware shall be less than 40 ps per C in the range of 20 to 30 C ambient room temperature. If necessary, active temperature control may

be used to meet the stability requirements. Figure 3 shows the two sample (Allan) deviation of the temperature (in units of C) in the room in which the measurement system will operate. For a given temperature sensitivity S, in seconds/ C, the corresponding Allan deviation of the measurement system due to temperature is $_y( ) = S \ _T( )/$

The stability of the measurement system shall meet the required levels independent of the number of clocks being measured. Of particular concern in this regard is the cross-talk noise between different channels in the measurement system. This noise tends to be most serious between adjacent channels, and the Contractor shall exercise great care in the design of the power-supply de-coupling filters and in isolating the grounds from the power-supply return lines.

Reliability

The system shall operate continuously for many years so a high degree of reliability is mandatory. It shall be designed for ease of repair. Given the performance and mechanical constraints outlined above, the system shall be as modular as possible and all of its modules shall be designed conservatively so that they operate at significantly less than their rated capacities. It is very desirable, though not mandatory, that the system be designed such that the measurement modules can be changed while the system is powered up.

DELIVERY AND TEST SCHEDULE

The first phase shall require the delivery of two fully functional systems with each containing at least 4 measurement channels at 5 and 100 MHz and the computer interface. Delivery shall be within 240 days after notice of award. With the exception of the driver software discussed above, NIST will be responsible for the software that will control the system and analyze the data. Contractor-supplied test programs may also be used, provided that they are fully documented and contain complete source-code listings. Test programs that do not satisfy these requirements may also be used for a preliminary evaluation, but data obtained using them shall not satisfy the testing requirements outlined here. These modules will be for the purposes of evaluating the stability and reliability of the systems. The testing phase will last on the order of 180 days.

DOCUMENTATION

Documentation shall include the full set of interface commands and a full set of block diagrams and schematics for the hardware. Source code for the driver software is highly desirable and is required if the driver will not be maintained by its supplier for a minimum of three years following delivery. Detailed information (regarding component selection or circuit design, for example) can be held by NIST as proprietary material if necessary, but nondisclosure agreements on system architecture or on the overall general principles of operation will not be accepted.